arXiv:1607.01489v2 [cond-mat.stat-mech] 27 Mar 2017

Estimating the speed-up of Adaptively Restrained Langevin Dynamics

Zofia Trstanova and Stephane Redon

Univ. Grenoble Alpes, LJK, F-38000 Grenoble, France; CNRS, LJK, F-38000 Grenoble, France; Inria, France

Abstract

We consider Adaptively Restrained Langevin dynamics, in which the kinetic energy function vanishes for small velocities. Properly parameterized, this dynamics makes it possible to reduce the computational complexity of updating inter-particle forces, and to accelerate the computation of ergodic averages of molecular simulations. In this paper, we analyze the influence of the method parameters on the total achievable speed-up. In particular, we estimate both the algorithmic speed-up, resulting from incremental force updates, and the influence of the change of the dynamics on the asymptotic variance. This allows us to propose a practical strategy for the parametrization of the method. We validate these theoretical results by representative numerical experiments.

Keywords:
Langevin dynamics, ergodic averages, variance, adaptive methods, computational complexity, sampling efficiency

1. Introduction

One fundamental aim of molecular simulations is the computation of macroscopic quantities, typically through averages of functions of the variables of the system with respect to a given probability measure μ, mostly distributed according to the Boltzmann-Gibbs density. This corresponds to a model of a conservative system in contact with a heat bath at a fixed temperature. Numerically, the high-dimensional averages with respect to μ

are often approximated as ergodic[1] averages over realizations of appropriate stochastic differential equations (SDEs) (see Balian (2006)).

In many applications, the dynamics is metastable, *i.e.* the system remains trapped for a very long time in some region of the configuration space. An acceleration of the sampling can be achieved by improving the exploration of the phase space (with variance reduction techniques such as importance sampling, see for instance Lelièvre et al. (2010); Lelièvre and Stoltz (2016)), increasing the time step by stabilizing the dynamics (see for instance Leimkuhler et al. (2016); Fathi and Stoltz (2015)), by changing the model as for example dissipative particle dynamics (see for instance Espanol and Warren (1995)), or the implementation (parallelization, reduction of the computational complexity (see for instance Bosson et al. (2012))). Many methods are based on a modification of the dynamics. Since, very often, the interest lies in computing of average properties, sampling can be unbiased to retrieve averages with respect to the canonical distribution. In order to increase the time step size used in the simulation, some methods consider modifying the kinetic energy based on changing the mass matrix (Bennett (1975); Plechac and Rousset (2010)). Another example is the Shadow-Hamiltonian Metropolis-Hastings method introduced by Izaguirre and Hampton (2004), which consists in integrating the Hamiltonian dynamics according to the Shadow-Hamiltonian, which is preserved by the numerical scheme.

In order to propose a fair comparison of sampling methods, three factors should be taken into account: the asymptotic variance of time averages, the maximal admissible time step size in the discretization and the computational effort.

In this work, we analyze the efficiency of a method based on a modified version of Langevin dynamics called "Adaptively Restrained Particles Simulations" (ARPS), first proposed in Artemova and Redon (2012). The main idea is to modify the kinetic energy function in order to freeze a number of particles at each time step and reduce the computational cost of updating inter-particles forces. In Artemova and Redon (2012), the kinetic energy is set to 0 when momenta are smaller than the restraining parameter $v_{\min}$, and is set to the standard, quadratic kinetic energy for momenta larger that the full-dynamics parameter $v_{\max}$, with $v_{\max} > v_{\min} \geq 0$. Thanks to this formu-

[1]In this article we use the term ergodic for the convergence in infinite time, *i.e.* the mathematical convergence (see Lelièvre and Stoltz (2016)).

lation, the computational complexity of the force update is reduced, because some particles do not move and hence, forces need not be updated. The associated gain can be quantified by an algorithmic speed-up factor $S_{\mathrm{a}} \geq 1$. On the other hand, since the dynamics is modified, the asymptotic variance of time averages σ^2_{AR} given by the Central Limit Theorem, differs from the asymptotic variance σ^2_{std} of the standard Langevin dynamics. Intuitively, iterates are a priori more correlated, which may translate into an increase of the statistical error. The actual speed-up of the method in terms of wall-clock time is therefore an interplay between the algorithmic speed-up and the variances. A rigorous mathematical analysis of the ergodicity properties of this method was provided in Redon et al. (2016). Moreover, a perturbative regime study of the asymptotic variance suggests a linear behavior of the variance with respect to the parameters of the dynamics in some limiting regime.

Since the method is parameterized by two constants, it is not obvious how to choose these parameters in order to achieve an optimal speed-up. Of course, the algorithmic speed-up depends on the percentage of restrained particles. The percentage of restrained particles is a non-linear function of the parameters, hence it is not trivial how to best choose their values. Our aim in this paper is to propose a strategy for choosing the parameters of the AR method.

The article is organized as follows: in Section 2.1 we make a brief overview of sampling using Langevin dynamics and we recall common strategies for its discretization. In Section 2.2 we recall the definition of AR-Langevin dynamics proposed in Artemova and Redon (2012) and the alternative definition of the AR-kinetic energy with better stability properties from Stoltz and Trstanova (2016). In Section 3 we give a definition of speed-up and we introduce a formula for the total speed-up with the AR approach. In the next two sections we analyze how this formula depends on the parametrization: in Section 4 we analyze the computational complexity of the method and we express the corresponding algorithmic speed-up. This part is followed by Section 5, in which we give a relation between the restraining parameters and the percentage of restrained particles, as well as an approach for obtaining the linear approximation of the variance with respect to the restraining parameters. By combining all the necessary parts, we propose a practical strategy for the parametrization of the method and we illustrate the theoretical results by numerical simulations in Section 6.

2. Modified Langevin dynamics

In this section we recall the Langevin dynamics and the modified Langevin dynamics.

2.1. Sampling from the canonical distribution using Langevin dynamics

We consider a system of N particles in a simulation box of the space dimension D with periodic boundary conditions. The total dimension of the system is $d = N \times D$. We denote by $p = (p_1, \dots, p_N) \in \mathbb{R}^d$ momenta of the particles and by $q = (q_1, \dots, q_N) \in \mathcal{D}$ their positions in the box $\mathcal{D} = (L\mathbb{T})^d$, where $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ is the one-dimensional unit torus. We denote $\mathcal{E} := \mathbb{R}^d \times \mathcal{D}$. The Hamiltonian of the system, which is the sum of the kinetic (K) and the potential energy (V), reads

$$H(q,p) = K(p) + V(q)\,.$$

Let us emphasize that we restrain ourselves to separable Hamiltonians.

The Langevin equations read:

$$\begin{cases} dq_t = \nabla K(p_t)\,dt, \\ dp_t = -\nabla V(q_t)\,dt - \gamma \nabla K(p_t)\,dt + \sqrt{\dfrac{2\gamma}{\beta}}\,dW_t\,, \end{cases} \tag{1}$$

where dW_t is a standard d-dimensional Wiener process, $\gamma > 0$ is the friction constant and $\beta = (k_{\mathrm{B}}T)^{-1} > 0$ is proportional to the inverse temperature. We refer the reader to Lelièvre and Stoltz (2016) for an overview of mathematical properties of this dynamics. The invariant distribution (the Boltzmann distribution) is simply obtained as

$$\mu(q,p) = Z^{-1}\mathrm{e}^{-\beta H(q,p)}, \quad Z = \int_{\mathcal{E}} \mathrm{e}^{-\beta H(q,p)} dq dp\,,$$

where Z is the normalization constant or the partition function. We use the notation μ_{std} for the case of the standard kinetic energy function, and μ_K for a general kinetic energy function. Langevin dynamics generate samples (p_t, q_t) from the Boltzmann distribution μ which are used in computation of macroscopic properties. These correspond to expected values with respect to μ and can be approximated by ergodic averages of the trajectories (p_t, q_t):

$$\lim_{t\to\infty} \hat{\varphi}_t = \mathbb{E}_\mu(\varphi) = \int_{\mathcal{E}} \varphi d\mu, \quad \hat{\varphi}_t = \frac{1}{t}\int_0^t \varphi(q_s, p_s) ds\,. \tag{2}$$

Note that for *any* well-defined[2] K, due to the separability of the Hamiltonian, the marginal distributions in the position variable remain unchanged, since only the momenta marginals of the distribution are influenced by the modification of the kinetic energy, *i.e.*

$$\mathbb{E}_{\mu_{\mathrm{std}}}\left[\varphi(q)\right] = \mathbb{E}_{\mu_K}\left[\varphi(q)\right] .$$

Since the dynamics (1) cannot be integrated exactly its solutions are approximated by numerical integration (Milstein and Tretyakov (2013)). Basically, there are two kinds of errors occurring in the estimation of $\mathbb{E}_\mu\left[\varphi\right]$ by $\hat{\varphi}_t$ from the numerical integration of equations (1): a statistical error, due to the finiteness of the time interval during which the sampling is performed; and a systematical error (bias) on the measure.

The statistical error for ergodic averages (2) is quantified by the Central Limit Theorem. The asymptotic variance associated with the estimator $\hat{\varphi}_t$ reads

$$\sigma^2 = \lim_{t\to\infty} t\mathbb{E}_\mu\left[\left(\Pi\varphi_t\right)^2\right] = 2\int_0^\infty \mathbb{E}\left[\Pi\varphi(q_s,p_s)\Pi\varphi(q_0,p_0)\right]ds, \quad \Pi\varphi := \varphi - \int_{\mathcal{E}} \varphi d\mu\,.$$

Similarly, for the discretized dynamics, with time step size $\Delta t = T/N_{\mathrm{iter}}$, we denote the estimator

$$\hat{\varphi}_{N_{\mathrm{iter}},\Delta \mathrm{t}} = \frac{1}{N_{\mathrm{iter}}}\sum_{i=0}^{N_{\mathrm{iter}}-1} \varphi(q^n,p^n).$$

If the discretized dynamics is geometrically ergodic with an invariant measure $\mu_{\Delta t}$, a Central Limit Theorem holds true and the variance of the discretized process is given by (see Meyn and Tweedie (1993); Lelièvre and Stoltz (2016))

$$\begin{aligned}
\sigma^2_{\Delta t} &= \lim_{N_{\mathrm{iter}}\to\infty} N_{\mathrm{iter}}\mathrm{Var}_{\mu_{\Delta \mathrm{t}}}\left(\hat{\varphi}_{\mathrm{N_{iter},\Delta t}}\right) \\
&= \mathbb{E}_{\mu_{\Delta t}}\left(\left[\Pi_{\Delta t}\varphi\right]^2\right) + 2\sum_{n=1}^{\infty}\mathbb{E}_{\mu_{\Delta t}}\left[\Pi_{\Delta t}\varphi(q^n,p^n)\Pi_{\Delta t}\varphi(q^0,p^0)\right] \qquad (3)\\
&= \int_{\mathcal{E}} \Pi_{\Delta t}\varphi\left[\left(2\left(\mathrm{Id}-P_{\Delta t}\right)^{-1} - \mathrm{Id}\right)\Pi_{\Delta t}\varphi\right]\, d\mu_{\Delta t}
\end{aligned}$$

[2]We assume that V belongs to $C^\infty(\mathcal{D},\mathbb{R})$, and $K \in C^\infty(\mathbb{R}^d,\mathbb{R})$ grows sufficiently fast at infinity in order to ensure that $\mathrm{e}^{-\beta K} \in L^1(\mathbb{R}^d)$.

where

$$\Pi_{\Delta t}\varphi := \varphi - \int_{\mathcal{E}} \varphi d\mu_{\Delta t}.$$

In other words, for N_{iter} simulation steps, the statistical error is of order $\varepsilon := \frac{\sigma_{\Delta t}}{\sqrt{N_{\text{iter}}}}$. The variance of the discretized process converges to the variance of the continuous process σ^2 as Δt tends to 0, i.e. $\Delta t \sigma^2_{\Delta t} \to \sigma^2$ as $\Delta t \to 0$ (see Lelièvre and Stoltz (2016)).

There are many possible ways to discretize (1), see for instance Leimkuhler et al. (2016); Mattingly et al. (2002); Kopec (2014) for a precise analysis of the properties of discretization schemes of the Langevin dynamics based on a splitting. A standard choice for the discretization of (1) is a numerical scheme of second order on the averages with respect to the time step size. It is possible to design higher order schemes, however they include at least double evaluation of the forces, which is not favorable due to the system size. Usually, the numerical schemes are constructed through a splitting of the generator of the Langevin dynamics $\mathcal{L} = A + B + O$ with

$$A := \nabla K(p) \cdot \nabla_q, \qquad B := -\nabla V(q) \cdot \nabla_p, \qquad O := -\gamma \nabla K(p) \cdot \nabla_p + \frac{\gamma}{\beta} \Delta_p \,.$$

For instance, the first order splitting (Lie-Trotter) gives the following scheme:

$$\begin{cases} \tilde{p} = p^n - \nabla V(q^n) \Delta t, \\ q^{n+1} = q^n + \nabla K(\tilde{p}) \Delta t, \\ p^{n+1} = \tilde{p} - \gamma \nabla K(\tilde{p}) \Delta t + \sqrt{\dfrac{2\gamma \Delta t}{\beta}} G^n \,, \end{cases} \tag{4}$$

where G^n are independent and identically distributed (i.i.d.) standard d-dimensional Gaussian random variables. This is the so-called BAO scheme. The name is motivated by the fact that the transition kernel reads $P^{\text{BAO}}_{\Delta t} \varphi = P^{\text{B}}_{\Delta t} P^{\text{A}}_{\Delta t} P^{\text{O}}_{\Delta t} \varphi$ where $P^A_{\Delta t}$ (respectively $P^B_{\Delta t}, P^O_{\Delta t}$) is the transition operator associated with the splitting step. We refer to Stoltz and Trstanova (2016) for a construction of second order discretization schemes in the case of a general kinetic energy.

Remark 2.1. *The Lie-Trotter and the Strang splitting each give six possible numerical schemes according to the order of the operators A,B,O (Leimkuhler et al. (2016)). Due to the high dimensionality of the system, the bottleneck of*

the computational complexity is the computation of the interactions between the particles, which must be done after each update of the positions (after action A). In addition, a non-negligible computational effort involves generation of random numbers in O. Therefore, schemes which include as few actions of A and O as possible should be preferred for a lower computational complexity.

2.2. AR-Langevin dynamics

In the usual setting, the kinetic energy function of system of N particles is a sum of kinetic energies of each particle, which are quadratic functions of momenta:

$$K_{\text{std}}(p) = \sum_{i=1}^{N} k(p_i)\,, \quad k(p_i) = \frac{p_i^2}{2m_i},$$

where $p_i \in \mathbb{R}^D$ is the momentum vector of the particle i with mass m_i.

AR-Langevin dynamics, proposed in Artemova and Redon (2012), is based on a modified kinetic energy function K that is defined as the sum of the kinetic energies of the individual particles, which are parametrized by two constants $v_{\max} > v_{\min} \geq 0$. In Artemova and Redon (2012), the AR-kinetic energy was designed such that it vanishes for values smaller than the restraining parameter $v_{\min}$, and are equal to the standard kinetic energy for values bigger than the full dynamics parameter $v_{\max}$. The main idea is that the derivative $\partial_{p_i} k$ of such function vanishes in this case too, and the position of the particle does not change between the two integration steps, i.e. $q_i^{n+1} = q_i^n, i \in I_R$, with I_R being the set of indices of the restrained particles (see also Equation (4) for the computation of q^{n+1}). The transition between the restraining and the full-dynamics region is performed with an interpolation spline, which ensures the regularity of the kinetic energy[3]. Still, the derivatives of k have large values in the transition region, which might cause numerical instabilities. However, the necessary condition for the particle to remain at the same position between two time steps is that the first derivative of the kinetic energy function vanishes for some values of momenta in every space direction. In Stoltz and Trstanova (2016) an alternative definition of the AR-kinetic energy function with this property was introduced, based on the vanishing velocities. The AR-kinetic energy is defined starting

[3]The choice was $k \in C^2(\mathbb{R}^D)$ in Artemova and Redon (2012).

from an interpolation of its first derivative, such that it vanishes around zero and takes values of p_i/m_i outside the restraining region (see Figure 2 for an illustration). The order of the interpolation spline can be chosen as high as necessary. However in this article we use only linear interpolation. The modified kinetic energy is then obtained by piecewise integration (see Figure 1 for an illustration):

$$K(p) = \sum_{i=1}^{d} k(p_i) \quad \text{where} \quad k(p_i) = \begin{cases} S_{v_{\min},v_{\max}} & \text{for} \quad \dfrac{|p_i|}{m_i} \leq v_{\min}, \\ s\left(p_i\right) & \text{for} \quad \dfrac{|p_i|}{m_i} \in [v_{\min}, v_{\max}], \\ \dfrac{p_i^2}{2m_i} & \text{for} \quad \dfrac{|p_i|}{m_i} \geq v_{\max}, \end{cases} \tag{5}$$

for $i = 1, \ldots d = N \times D$, and where s is the integrated interpolation spline and $S_{v_{\min},v_{\max}}$ is an integration constant, which corresponds to the minimal kinetic energy value of the particle[4]. Note that the total kinetic energy is the sum of individual kinetic energies of each particle in every space direction.

AR dynamics accelerate sampling by exploiting information about the kinetic energy of particles. More precisely, a particle is called *restrained* if it has the absolute value of each component of its momentum smaller than the restraining threshold $v_{\min}$. All other particles are defined as *active* particles. Note that, during the simulation, particles are switching between these states. Also, the average occupation of the active or restrained state only depends on the restraining parameters $v_{\min}$ and $v_{\max}$.

Since the momenta of individual particles are independent from each other under the canonical measure, the parameters $v_{\min}$ and $v_{\max}$ could in fact be different for each particle, to either focus calculations on a specific part of the particle system or to adjust the scaling of parameters according to the mass of the particle.

3. Estimating the speed-up

In this section we introduce a framework for the complexity analysis of the AR dynamics in the case of pair-wise interactions, which are the most

[4]We would like to emphasize that this constant does not appear in the Langevin equations, only in the momenta marginal of the invariant measure.

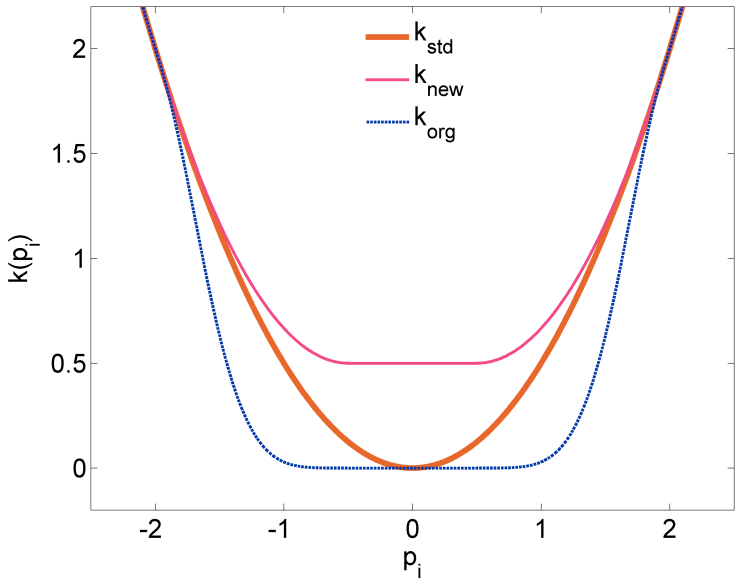

Figure 1: *Comparison of the AR-kinetic energy functions for one particle ($p = p_0, N = 1$) (Stoltz and Trstanova (2016))*: k_{org} is the original definition from Artemova and Redon (2012), k_{new} corresponds to Equation (5) and k_{std} is the standard one.

common interactions in numerous applications. Note that the discussion below can be easily generalized to interactions present in classical force-fields. The force acting on each particle i is a sum of interactions with all other particles.

The information about the state of the particle allows us to lower the computational cost of the computation of pair-wise interactions between the particles. We consider the potential

$$V(q) = \sum_{\substack{j=1 \\ i \neq j}}^{N} v(r_{ij})$$

and the force acting on the particle i which is given by

$$f_i(q) := -\partial_{q_i} V(q) = -\sum_{\substack{j=1 \\ i \neq j}}^{N} v'(r_{ij}) \frac{q_i - q_j}{r_{ij}}, \quad r_{ij} = |q_i - q_j| \,. \tag{6}$$

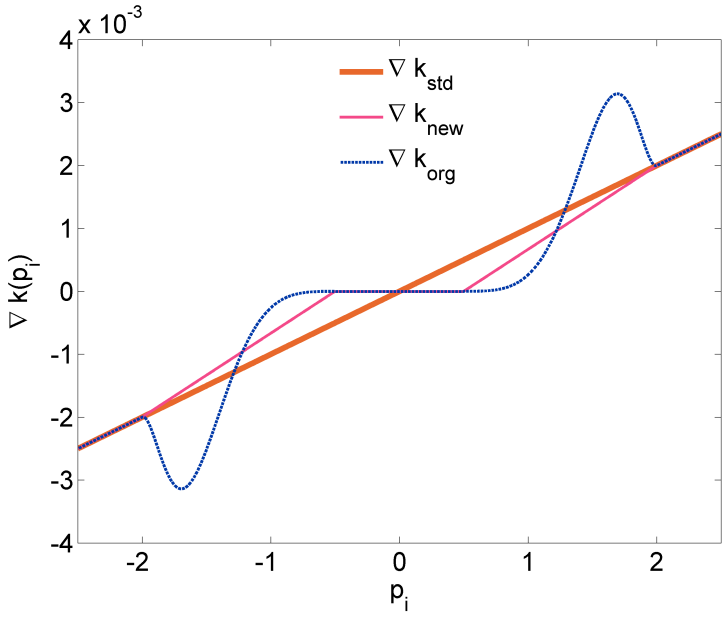

Figure 2: *Gradient interpolation of the kinetic energy (k_{new}) versus function interpolation (k_{org})*, see also Figure 1. Both gradients are vanishing around zero, which allows to freeze the positions of the particles. However, the gradient interpolation has been proven to be more numerically stable.

The change of the force between two time steps only depends on active particles that have moved since the last time step with respect to this particle. This allows us to avoid the computation of pair-wise interactions between restrained particles, hence lower the computational complexity. In order to quantify the computational cost of the force update, we define the force function $\psi : \mathbb{R} \to \mathbb{R}$ such that $\psi := v'$. Then the *computational cost of the force update* is defined as the number of times the force function ψ is called. The speed-up of AR dynamics, due to the decreasing of the computational complexity in the force update, with respect to the non-adaptive method which updates all interactions, defines *the algorithmic speed-up*. Since the computational complexity depends on the ratio of restrained particles, which is a quantity that varies at each time step, we consider averages over the whole simulation. More precisely, we denote by $C_{\text{AR},n}$ the computational cost of the force update in the AR-method at time step n and by $C_{\text{std},n}$ the computational cost of a standard, non-adaptive method. We denote by $N_{\text{iter}} = T/\Delta t$ the number of time steps in the simulation. Then *the algorithmic speed-up*

S_{a} is the ratio of the average computational cost $\widehat{C}_{\mathrm{std}} := \mathbb{E}_{\mu_{\mathrm{std}}}[C_{\mathrm{std}}]$ in the standard method and the average computational cost $\widehat{C}_{\mathrm{AR}} := \mathbb{E}_{\mu_{\mathrm{AR}}}[C_{\mathrm{AR}}]$ in the AR method:

$$S_{\mathrm{a}} := \frac{\lim_{N_{\mathrm{iter}}\to\infty} \frac{1}{N_{\mathrm{iter}}} \sum_{n=0}^{N_{\mathrm{iter}}} C_{\mathrm{std},n}}{\lim_{N_{\mathrm{iter}}\to\infty} \frac{1}{N_{\mathrm{iter}}} \sum_{n=0}^{N_{\mathrm{iter}}} C_{\mathrm{AR},n}} = \frac{\widehat{C}_{\mathrm{std}}}{\widehat{C}_{\mathrm{AR}}}\,. \tag{7}$$

Note that the computational complexities in both cases are bounded functions of the number of particles and, due to the ergodicity of the dynamics, which was proved by Redon et al. (2016), the averages in (7) almost surely converge.

However, the important point is the reduction of the error for a given wall-clock duration. We focus here on the statistical error, which is often the dominant source of errors. In order to express the total speed-up with respect to the standard method, we need to consider not only the algorithmic speed-up, but also the modification of the asymptotic variance which depends on the concrete choice of the kinetic energy (see expression (3)). We define *the total speed-up* S_{total} as a ratio of the wall-clock time, which is needed by using the AR-method in order to achieve some statistical precision, and the wall-clock time needed for reaching the same precision by the standard method:

$$S_{\mathrm{total}} := \frac{T^{\mathrm{wlck}}_{\mathrm{std}}}{T^{\mathrm{wlck}}_{\mathrm{AR}}}\,. \tag{8}$$

Recall that, for an observable φ, we denoted by $\sigma^2_{\Delta t}$ the asymptotic variance of the sampling from the discretized dynamics and by σ^2 the asymptotic variance of the continuous dynamics. From the Central Limit Theorem, the statistical error at time T is given by

$$\hat{\varphi}_T = \mathbb{E}_\mu(\varphi) + \varepsilon_T \mathcal{G},$$

where ε_T is of order $\frac{\sigma}{\sqrt{T}}$ and $\mathcal{G} \sim \mathcal{N}(0,1)$. Hence the number of time steps $N_{\mathrm{iter}} = T/\Delta t$ needed in order to have a statistical error of order ε_T is

$$N_{\mathrm{iter}} = \frac{\sigma^2_{\Delta t}}{\varepsilon_T^2}\,.$$

The corresponding wall-clock time is therefore obtained by considering the average cost $\widehat{C}$ as

$$T^{\text{wlck}} = N_{\text{iter}}\widehat{C}\,.$$

Taking into account that $\Delta t\sigma^2_{\Delta t} \sim \sigma^2$ (for time steps small enough, recall Section 2.1), the total speed-up S_{total} defined in (8) can be expressed as

$$S_{\text{total}} = \frac{\widehat{C}_{\text{std}}}{\widehat{C}_{\text{AR}}}\frac{\sigma^2_{\text{std},\Delta t}}{\sigma^2_{\text{AR},\Delta t}} = S_{\text{a}}\frac{\sigma^2_{\text{std},\Delta t}}{\sigma^2_{\text{AR},\Delta t}} \approx S_{\text{a}}\frac{\sigma^2_{\text{std}}}{\sigma^2_{\text{AR}}}\frac{\Delta t_{\text{AR}}}{\Delta t_{\text{std}}}\,. \tag{9}$$

The last two terms in (9) become equal for small values of Δt and it is therefore sufficient to study the variance of the continuous process σ^2_{std} and σ^2_{AR}. As we have already mentioned, the choice of the modified kinetic energy should not change the stability properties of the standard dynamics. This would otherwise require us to choose a smaller time step size Δt_{AR}, which would lead to a smaller total speed-up S_{total}. Unfortunately, this is the case of the kinetic energy defined in Artemova and Redon (2012). Still, the stability can be significantly improved by using the kinetic energy given by (5) instead. In this case, the stability properties become comparable to the ones of the standard dynamics (Stoltz and Trstanova (2016)). We therefore assume in this work $\Delta t_{\text{std}} = \Delta t_{\text{AR}}$.

The computation in (9) shows the trade-off between the algorithmic speed-up and the change in variance. Both the algorithmic speed-up S_{a} and the AR variance σ^2_{AR} depend on the parameters of the AR dynamics. As already showed in Artemova and Redon (2012), in some applications, the restraining parameters can be chosen such that the total speed-up satisfies $S_{\text{total}} > 1$. Therefore, there are systems for which this method can be efficient, even though this might be counter-intuitive since one could suggest that in order to accelerate the sampling, the system should move "faster" and not be restrained. Note however that the wall-clock duration of the force update step depends on the implementation and on the complexity of the evaluation of ψ. Hence, the same physical model with variance σ^2, can have various algorithmic speed-ups S_{a}. Finally, an interesting observation is, that due to the separability of the Hamiltonian, the algorithmic speed-up does not depend on the potential.

4. Algorithmic speed-up

The goal of this section is to propose a methodology to analyze the algorithmic speed-up S_{a} (defined in (7)) of AR dynamics as a function of the

percentage of restrained particles. We first describe the adaptive algorithm for computing forces, and we estimate the corresponding computational cost. In the last part, we also consider the effort for updating neighbor lists used for updating of short-ranged interactions and we obtain an estimation of the algorithmic speed-up per time step.

4.1. Description of the AR force update algorithm

For simplicity, we consider a system of N particles where only pair-wise interactions take place. In AR dynamics, this sum can be split into three kinds of interactions depending on the state of the two interacting particles: active-active, active-restrained and restrained-restrained. We define the set of indices of active particles I_{A} and restrained particles I_{R}. Then, sum (6) can be re-written as

$$f_i = \sum_{\substack{j \in I_{\mathrm{A}} \\ j \neq i}} f_{ij} + \sum_{\substack{j \in I_{\mathrm{R}} \\ j \neq i}} f_{ij} \,. \tag{10}$$

The force acting on particle i in the next time step $n+1$ can be formally obtained using the old position q^n:

$$f_i^{\mathrm{new}} = f_i^{\mathrm{old}} + \left(f_i^{\mathrm{new}} - f_i^{\mathrm{old}} \right) \,, \quad f_i^{\mathrm{old}} = \sum_j f_{ij}(q^n), \quad f_i^{\mathrm{new}} = \sum_j f_{ij}(q^{n+1}) \,. \tag{11}$$

Since, for the set of restrained particles, positions have not changed since the previous time step, one can easily see that

$$\forall i \in I_R, \quad \sum_{\substack{j \in I_{\mathrm{R}} \\ j \neq i}} f_{ij}^{\mathrm{new}} - f_{ij}^{\mathrm{old}} = 0 \,.$$

The computation in (11) is therefore reduced to subtracting the old and adding the new active-restrained and active-active interactions. This simple remark provides in fact a key point for the reduced complexity of the AR algorithm.

In a standard simulation, when taking into account Newton's third law $f_{ij} = -f_{ji}$, the computational cost of pair-wise interactions is $\frac{N(N-1)}{2}$. The resulting quadratic complexity in the number of particles is not favorable due to the system size, and therefore neighbor lists are usually introduced (see Allen and Tildesley (1989); Frenkel and Smit (2002)). For a comparison of various approaches for neighbor list methods we refer the reader to

Artemova et al. (2011). Neighbor lists can be used in systems where forces vanish after a certain cut-off distance, so that each particle only interacts with a relatively limited number of neighbors. For simplicity, we consider a homogenous system where we assume that the number of neighbors C of a particle is the same for each particle. Taking into account that, for each pair (i, j), we may only compute the force f_{ij} and deduce f_{ji} thanks to Newton's third law, the number of interactions reduces to $\frac{NC}{2}$.

The BAO discretization scheme (4) can be formalized in the following way:

Input: Initial conditions p^0, q^0
Output: p^n, q^n
for *each time step and each particle* **do**
    B: Update momenta ;
    A: Update positions ;
        Update neighbor-lists ;
        Update forces ;
    O: Update momenta in fluctuation-dissipation part (FD) ;
end

Algorithm 1: Algorithm for (4) in the case of the standard dynamics.

In AR dynamics, the information about which particles are going to move after the position update is already available after updating the momenta B, since the kinetic energy will not change anymore. The algorithm above may thus be modified in the following way:

Input: Initial conditions p^0, q^0
Output: p^n, q^n
For each particle i, initialize force f_i;
for *each time step and each particle* **do**
  B: Update momenta;
    Create lists of active and restrained particles;
    Subtract active-active and active-restrained interactions;
  A: Update positions;
    Update neighbor-lists;
    Add active-active and active-restrained interactions;
  O: Update momenta in FD;
end

Algorithm 2: Algorithm for (4) in AR dynamics using adaptive forces updates.

Updating neighbor lists normally consists in re-assigning each particle to a specific grid cell (in our implementation we used a combination of Verlet lists and linked-cell lists (Frenkel and Smit (2002))). In AR dynamics, restrained particles do not have to be re-assigned, and neighbor lists may be updated more efficiently. More precisely, the complexity of updating the neighbor lists goes from $O(N)$ the number of particles, to $O(K)$, where K is the number of active particles.

Note that the force function ψ is called in both AR force updates (subtract and add steps), since we need to evaluate forces for positions at the previous time step. It would be possible to avoid updating forces twice by saving all pairwise forces, but this may result in a quadratic space complexity. We will not analyze this case, although it would result in a larger algorithmic speed-up and lead to less restrictive conditions on the efficiency of AR dynamics.

Note that there is a slight overhead due to computing the AR kinetic energy functions ∇K, which is more complicated than in the standard case. Still, this involves $O(N)$ additional operations, and can be neglected compared to the cost of the force update in practical applications. Furthermore, the overhead mostly comes from the transition regime since ∇K vanishes for restrained particles and becomes $M^{-1}p$ for the full-dynamics state.

Note that a similar strategy for incremental force update may be applied to other splitting schemes of the modified Langevin equations. However, the status of a particle (active, in transition or restrained) depends on the state of the momenta before the position update, and hence this status should not

be destroyed by updating momenta between two position updates. Using the same notation as in Section 2.1 and Leimkuhler et al. (2016), this implies that the second order splittings BAOAB and OABAO are not directly suited for a modification by the AR dynamics algorithm, since between the two position updates A, the momentum changes by either O or B step. On the other hand, ABOBA, BOAOB, OBABO and AOBOA can be used, since the lists of active particles can be created before the position update A and hence active-active and active-restrained interactions can be subtracted and added after position update A.

4.2. Complexity analysis

At each time step, the computational cost of the force update depends on the percentage of restrained particles. Let us denote the number of active particles by $K = \alpha N$, where $\alpha \in [0, 1]$ is the average ratio of active particles. The number of restrained particles is then $N - K$. We are going to formalize the computational complexity of the force update as a function of the ratio of restrained particles, denoted by $\rho := 1 - \alpha$.

We recall that we have considered the average computational cost over the whole trajectory in equation (7), since the instantaneous computational cost may vary at each time step. Because, in the algorithm analyzed in this paper, we add and subtract pairwise forces, the computational complexity of the force update in an AR simulation is lower than a regular force update if and only if a sufficient number of particles is restrained. We are thus going to analyze which conditions on the number of restrained particles are sufficient to obtain a speed-up larger than one, when a standard simulation has a linear or quadratic complexity[5]. This analysis can be extended to other force update algorithms.

4.2.1. Quadratic complexity

Let us first consider a standard (non-adaptive) simulation with a quadratic-complexity force update algorithm, i.e. when no neighbor-lists are used. In this case, the number of interactions computed at every time step is $C_{\text{std}} := \frac{N(N-1)}{2}$. In AR dynamics, we do not need to recompute interactions between restrained particles, hence we only update interactions involving active particles, either with other active particles, or with restrained particles.

[5]The quadratic complexity corresponds to bonded interactions and the linear complexity to non-bonded, in which case the neighbor lists can be applied.

As a result, the computational cost for the AR force update is[6]:

$$C_{\mathrm{AR}} := 2\left(\frac{K(K-1)}{2} + K(N-K)\right) = (2N - \alpha N - 1)\, N\alpha\,.$$

and $C_{\mathrm{std}} > C_{\mathrm{AR}}$ is satisfied for

$$\alpha < 0.29 \tag{12}$$

and $N > \frac{1-2\alpha}{2\alpha^2-4\alpha+1}$. The inferior bound on the number of particles is not a restrictive condition for molecular dynamics, where the number of particles is usually much bigger. (For example, for $\alpha = 0.28$, the number of particles N must be larger than 12.) More importantly, this implies that at least 71% of particles must be restrained in order for this force update algorithm to be beneficial, in which case the algorithmic speed-up is:

$$S_{\mathrm{a},1} = \frac{C_{\mathrm{std}}}{C_{\mathrm{AR}}} = \frac{N-1}{2(N\rho + N - 1)(1-\rho)}\,.$$

When the number of particles tends to infinity, this becomes

$$S^{\infty}_{\mathrm{a},1} = \lim_{N\to\infty} S_{\mathrm{a}}(N) = \frac{1}{2(\rho+1)(1-\rho)}\,. \tag{13}$$

Note that if the double computation of forces can be avoided (for example by storing previous pairwise forces), the complexity becomes

$$C_{\mathrm{AR},2} := \frac{C_{\mathrm{AR}}}{2}\,,$$

so that $C_{\mathrm{AR},2} > C_{\mathrm{std}}$ is achieved for any $\alpha < 1$ and $N > \frac{1}{1-\alpha}$, resulting in the following speed-up:

$$S_{\mathrm{a},2} = \frac{N-1}{(N\rho + N - 1)\,(1-\rho)}, \quad S^{\infty}_{\mathrm{a},2} = \frac{1}{(\rho+1)(1-\rho)}\,.$$

[6]The factor of 2 comes from the need to subtract old forces (with previous positions) and add new forces (with current positions).

4.2.2. Linear complexity

Let us now consider the (much more frequent) case where the complexity of the force update is linear, e.g. when forces become sufficiently small after a given cutoff distance r_{cut}, and neighbor lists may be used to efficiently determine which particles are interacting. The reference complexity is therefore $C_{\text{std,NL}} = \frac{NC}{2}$, where C is the (average) number of neighbors. The algorithm for the adaptive force update is as follows: for all active particles compute interactions with their neighbors, and between the active neighbors use $f_{ij} = -f_{ji}$. The total number of interactions to be updated in the AR dynamics algorithm is then:

$$C_{\text{AR,NL}} := 2\left(\sum_{i=1}^{K}\sum_{j\in L_A(i)} 1 + \sum_{i=1}^{K}\sum_{j\in L_R(i)} 1\right) = 2\left(\frac{KC_A}{2} + KC_R\right) = \left(1-\frac{\alpha}{2}\right)\alpha CN\,, \tag{14}$$

where the set $L_A(i) \subset I_A$ contains the indices of the active neighbors of the particle i, $L_R(i) \subset I_R$ contains the indices of the restrained neighbors, $C_A = \alpha C$ and $C_R = \rho C$. The necessary condition for $C_{\text{std,NL}} > C_{\text{AR,NL}}$ is then

$$\alpha < 0.293\,. \tag{15}$$

Note that this condition does not depend on N, nor on C. The AR dynamics algorithm is more efficient in number of operations for forces update if and only if the percentage of restrained particles is bigger than $70,7\%$. The algorithmic speed-up is hence

$$S_{\text{NL}} := \frac{C_{\text{std}}}{C_{\text{AR,NL}}} = \frac{1}{2\alpha(2-\alpha)} = \frac{1}{2(1-\rho^2)}\,. \tag{16}$$

Again, avoiding the double re-computation of force components from the old positions for the active particles, removes a factor of 2 from $C_{\text{AR,NL}}$ and the computational cost becomes $C_{\text{AR,NL}} = \left(\frac{1}{2}-\alpha\right)\alpha CN$, which implies an unconditional algorithmic speed-up $S_{\text{a}} = (1-\rho^2)^{-1}$.

An important conclusion is that an incremental force update is computationally beneficial if the percentage of restrained particles is larger than a threshold $\mathcal{R}$. We may thus modify Algorithm 2 as follows:

Input: Initial conditions p^0, q^0. **Output**: p^n, q^n
For each particle i, initialize the force f_i;
for *each time step and each particle* **do**
  B: Update momenta;
    Create lists of active and restrained particles;
  if $\rho > \mathcal{R}$ **then**
    Subtract active-active and active-restrained interactions;
    A: Update positions;
    Add active-active and active-restrained interactions;
  else
    A: Update positions;
    Update forces with the standard approach;
  end
    Update neighbor-lists;
  O: Update momenta in FD;
end

Algorithm 3: Improvement of Algorithm 2 by using the condition on the ratio of restrained particles ρ given by the constant $\mathcal{R} \geq 0$.

Finally, we consider the case where the neighbor lists are updated at each time step[7]. This is not usually done in practical applications, where neighbor lists are updated only after a certain time period which can be computed from the maximal velocity of the particles. The cost per time step then extends in re-assigning N particles into the grid, which gives order of $NC/2 + N$ operations. In the AR simulation, only active particles need to be re-assigned into the grid. Therefore, the cost per time step computed in (14) becomes

$$C_{\mathrm{AR,NL}} := \left(1 - \frac{\alpha}{2}\right) \alpha C N + \alpha N \,.$$

Assuming that there are C neighbors in average, the resulting speed-up is:

$$S_{\mathrm{a}} = \frac{C + 2}{2(1-\rho)(C\rho + C + 1)} \,. \tag{17}$$

[7]Note that this can be easily modified in order to express the update of neighbor-lists every time period T.

5. Total speed-up

As explained above, the total speed-up (7) reachable by AR dynamics when estimating observables depends on both the computational complexity of the force update, and the variance of the AR dynamics.

In this section, we first analyze how the percentage of restrained particles depends on the restraining parameters $v_{\min}$ and $v_{\max}$. Then, we approximate the variance of the AR dynamics by a linear function. Combining both, we finally express the total speed-up as a function of $v_{\min}$ and $v_{\max}$.

5.1. Percentage of restrained particles

The percentage of restrained particles can be computed from the average occupation of the restrained state of each particle. In other words, it is the probability that the momenta of one particle belong to the restrained region of phase space. For the AR kinetic energy function (5), the average occupation of the restrained state $R(v_{\min}, v_{\max})$ of particle i with parameters $v_{\min}$ and $v_{\max}$ is the expected value of the indicator function of the absolute values of all momenta components of one particle being smaller than the restraining parameter $v_{\min}$.

We denote by $\mu_{v_{\min},v_{\max}}$ the invariant measure which corresponds to the AR kinetic energy function with parameters $v_{\min}$ and $v_{\max}$ and we compute

$$R(v_{\min}, v_{\max}) = \int_{\mathbb{R}^d} \mathbb{1}_{\left\{\frac{|p_i|}{m_i} \leq v_{\min}\right\}} \mu_{v_{\min},v_{\max}} = \frac{(2v_{\min}m_i)^D \exp\left(-\beta D S_{v_{\min},v_{\max}}\right)}{Z_p(v_{\min}, v_{\max})}\,, \tag{18}$$

where the momenta normalization constant of the particle i is simply $Z_p = z^D$, with

$$\begin{aligned} z(v_{\min}, v_{\max}) &= \int_{\left\{\frac{|p_i|}{m_i} \leq v_{\min}\right\}} \mathrm{e}^{-\beta S_{v_{\min},v_{\max}}}\, dp + \int_{\left\{\frac{|p_i|}{m_i} \geq v_{\max}\right\}} \mathrm{e}^{-\beta \frac{p_i^2}{2m_i}}\, dp \\ &\quad + \int_{\left\{\frac{|p_i|}{m_i} \in [v_{\min}, v_{\max}]\right\}} \mathrm{e}^{-\beta s(p_i)} dp \\ &= (2v_{\min}m_i)\, \mathrm{e}^{-\beta S_{v_{\min},v_{\max}}} + \sqrt{2\pi m_i \beta^{-1}} - 2\int_0^{m_i v_{\max}} \mathrm{e}^{-\beta \frac{p^2}{2m_i}} dp \\ &\quad + \int_{m_i v_{\min}}^{m_i v_{\max}} \mathrm{e}^{-\beta s(p_i)} dp + \int_{-m_i v_{\max}}^{-m_i v_{\min}} \mathrm{e}^{-\beta s(p_i)} dp\,. \end{aligned} \tag{19}$$

Note that in the standard dynamics $Z_p = (2\pi m_i/\beta)^{3/2}$.

Finally, considering a system consisting of particles with various restraining parameters $v^i_{\min}$ and $v^i_{\max}$, the total average percentage of restrained particles can be computed as an average over the individual values $R(v^i_{\min}, v^i_{\max})$ of each particle. Denoting by $N_{v_{\min},v_{\max}}$ the number of particles with parameters $v_{\min}$ and $v_{\max}$ and by $\mathcal{N}$ the set of all chosen pairs $(v_{\min}, v_{\max})$, the total average percentage of restrained particles[8] $R_{\text{total}} \in [0,1]$ is given by

$$R_{\text{total}} = \frac{1}{N} \sum_{(v_{\min},v_{\max})\in\mathcal{N}} N_{v_{\min},v_{\max}} R(v_{\min}, v_{\max})\,. \tag{20}$$

For example, the percentage of restrained particles for a system consisting of a dimer that follows standard dynamics and that is surrounded by solvent particles following AR dynamics with non-zero parameters $v_{\min}$ and $v_{\max}$ is:

$$R^{\text{DS}}_{\text{total}} = \frac{N_{\text{Solv}}}{N_{\text{total}}} R_{\text{Solv}}(v_{\min}, v_{\max})\,,$$

since, in standard dynamics, the average occupation of the restrained state is zero and $R_{\text{Dimer}} = 0$.

In conclusion, the algorithmic speed-up S_{a} can be estimated using the computational complexity of the algorithm (see Section 4) with the speed-up being a function of $\rho = R_{\text{total}}$.

Figure 3 shows, for K defined in (5), the average occupation of the restrained state $R(v_{\min}, v_{\max})$ as a function of $v_{\max}$ for various $\delta \in [0.5, 0.98]$ such that $v_{\max} = \delta v_{\min}$ in dimension three. We depicted also the value 70% of restrained particles which corresponds to the necessary condition for $S_{\text{a}} > 1$ (given by (12) or (15)). We observe on this figure that the bigger δ, the bigger average occupation of the restrained state. Figure 4 shows the dependence of $R(v_{\min}, v_{\max})$ on the temperature. This suggests that the restraining parameters should be scaled with respect to the temperature $k_B T$.

Finally, Figure 5 shows $R_{\text{total}}(v_{\min}, v_{\max})$ as a function of both parameters. Note that the highest value of percentage of restrained particles is located close to the diagonal, *i.e.* when the gap between the parameters $v_{\min}$ and $v_{\max}$ is small.

[8]Note that this corresponds to the notation $\rho = R_{\text{total}}$ in Section 4.

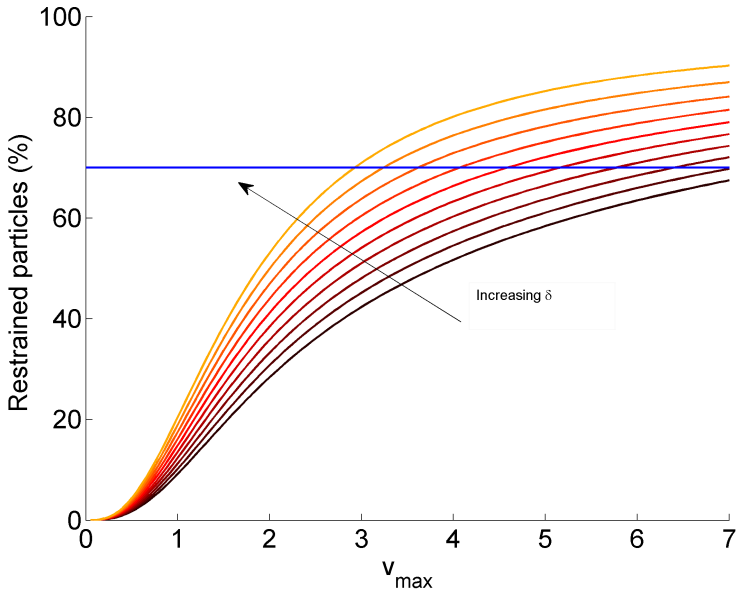

Figure 3: *Average occupation of the restrained state with respect to the parameter ratio.* We computed $R_{\text{total}} = R(v_{\min}, v_{\max})$ for one particle in 3D according to (18) for various $v_{\max}$ and various values of the parameter ratio $\delta \in [0.7, 0.98]$ (black to orange or bottom to top lines) such that $v_{\min} = \delta v_{\max}$. The blue line is the value 70% of restrained particles which corresponds to the necessary condition for $S_{\text{a}} > 1$.

5.2. Linear approximation of the variance

In Redon et al. (2016), it was proved that there exists a regime in which the variance from the AR dynamics simulations can be approximated by a linear function of the restraining parameters (see (Redon et al., 2016, Proposition 4.3)): there exists $v^*_{\max}$ small enough such that for $v_{\max} < v^*_{\max}$ there exist constants $c_1, c_2 \in \mathbb{R}$ such that for $v_{\min} \in [0, v_{\max}]$

$$\sigma^2_{\text{AR}}(v_{\min} + \zeta, v_{\max} + \eta) = \sigma^2_{\text{AR}}(v_{\min}, v_{\max}) + c_1\zeta + c_2\eta + \text{O}\left(\zeta^2 + \eta^2\right). \quad (21)$$

The total speed-up of the wall-clock time needed to reach a certain statistical precision (9) can hence be expressed in terms of the restraining parameters using (21) as

$$\begin{aligned} S_{\text{total}} &\approx S_{\text{a}}(v_{\min}, v_{\max}) \frac{\sigma^2(0,0)}{\sigma^2(0,0) + c_1 v_{\min} + c_2 v_{\max}} \\ &= S_{\text{a}}(v_{\min}, v_{\max}) \frac{1}{1 + \frac{c_1}{\sigma^2(0,0)} v_{\min} + \frac{c_2}{\sigma^2(0,0)} v_{\max}}. \end{aligned} \quad (22)$$

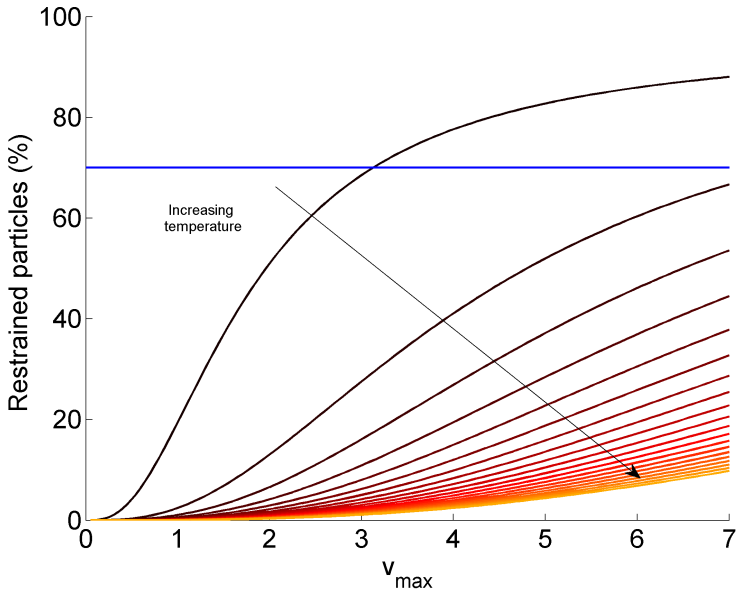

Figure 4: *Average occupation of the restrained state with respect to the temperature.* We computed $R_{\text{total}} = R(v_{\min}, v_{\max})$ of a particle in 3D according to (18) for $v_{\max}$ and $v_{\min} = 0.95 v_{\max}$ and for various temperatures $k_B T = [1, 100]$ (black to orange or top to bottom lines).

The gap between the restraining parameters $v_{\min}$ and $v_{\max}$ should be big enough to ensure a smooth transition between the full and the restrained dynamics and prevent numerical instabilities. Note, however, that in the numerical experiments performed in (Redon et al., 2016, Section 5.1), where the variance was computed for a simple 1D system, it was shown that the relative increase of the variance with respect to the full-dynamics parameter $v_{\max}$ is more significant than with respect to the restraining parameter $v_{\min}$. This result is not surprising, since the gap between the parameters smooths out the dynamics, which translates into an increase of correlations. This implies that the optimal strategy is to choose the gap between the parameters as small as possible while still maintaining the numerical stability and keeping the systematical error sufficiently low (i.e. the error on the computed averages, arising from the fact that $\mu \neq \mu_{\Delta t}$ (Leimkuhler et al. (2016))). At the same time, the restraining parameters should give the desired percentage of restrained particles R_{total}. For example, in the case when $\delta = v_{\min}/v_{\max} = 0.98$, the relative derivative of the restrained energy of one

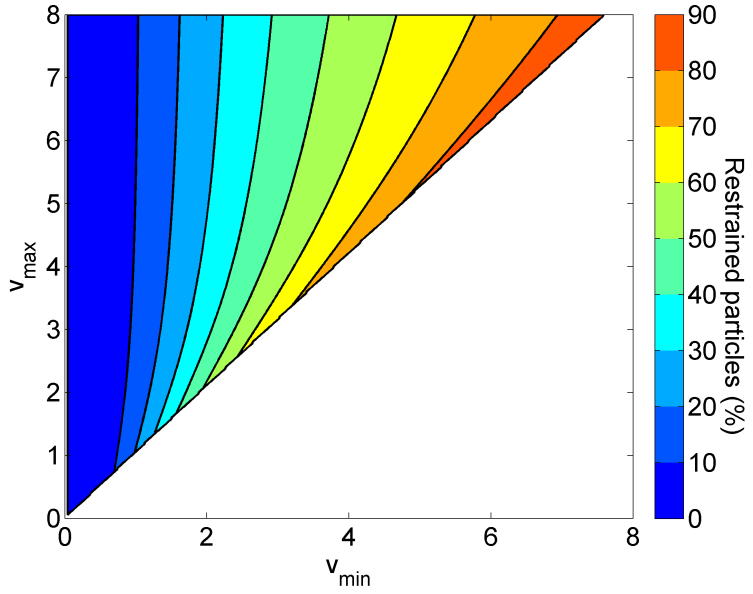

Figure 5: *Percentage of restrained particles over* $v_{\min}$ *and* $v_{\max}$. We considered the system of a dimer surrounded by a solvent (64 particles) in 3D with the solvent dynamics being governed by the AR kinetic energy function. We computed the percentage of restrained particles for various values $0 \leq v_{\min} \leq 0.95 v_{\max}$.

particle $R(v_{\min}, v_{\max})$ with respect to $v_{\max}$, almost vanishes after the value $v_{\max} = 5$. Hence this is a critical value after which the growth of function $R(v_{\min}, v_{\max})$ slows down (see again Figure 3). Having in mind that the variance locally increases with respect to $v_{\max}$, this implies that, in this region, the efficiency of the algorithmic speed-up does not grow fast enough with increasing $v_{\max}$, while the variance might be still growing. In this case, either the gap δ should be chosen smaller, or one must ensure that the variance does not grow too fast, in order to compensate the variance growth with the algorithmic speed-up.

It is easy to estimate the algorithmic speed-up S_{a}. The problematic part is to estimate the sensitivity of the variance of a given observable with respect to the modification by the restrained dynamics, i.e. the estimation of c_1 and c_2 in (22). This can be done by a linear interpolation in the pre-processing part, which should involve at least three AR dynamics simulations in order to estimate the constants c_1 and c_2. Finally, (22) allows to have a complete expression for the total speed-up as a function of the parameters $v_{\min}$ and

$v_{\max}$. Choosing $\delta = v_{\min}/v_{\max}$ as small as possible, one can find the optimal $v_{\max}$ which produces the highest total speed-up (see Section 6 for a numerical example).

We thus propose the following guidelines to estimate the total speed-up S_{total} with respect to the parameters $v_{\min}$ and $v_{\max}$:

1. Choose the order (scale) of the restraining parameters $v_{\min}$ and $v_{\max}$ for each particle according to its mass, its role in the system and the temperature $k_B T$.
2. Choose the minimal gap δ between $v_{\min}$ and $v_{\max}$ with respect to the numerical stability.
3. Compute the percentage of restrained particles R_{total} according to (18), (19) and (20).
4. Compute the algorithmic speed-up S_{a} according to the implementation algorithm according to Section 4.
5. Estimate the linear approximation of the variance $\sigma^2_{\text{AR}}(\delta v_{\max}, v_{\max})$ for the observable A.
6. Find the optimal value of $v_{\max}$ (with $v_{\min} = \delta v_{\max}$) by maximizing $S_{\text{total}}(v_{\min}, v_{\max})$.

6. Numerical illustration

In order to illustrate the theoretical results from the previous section, we consider a system of $N = 64$ particles consisting of a dimer (q_1, q_2) surrounded by 62 solvent particles $(q_3, \ldots, q_N)$ in space dimension $D = 3$ (see Figure 6 for an illustration of this system). This model is representative of many applications in biology, chemistry or physics, where the macroscopic property only depends on a small part of the simulated system. For example, in simulation of a protein in solvent the interest lies in computation macroscopic properties of the protein (see for instance Artemova and Redon (2012)). The validation of this method for real-world problems is still needed and this work has already started by the ARPS method being implemented into LAMMPS (see Singh and Redon (2016)).

We use periodic boundary conditions with box-length such that the density is 0.4. We consider reduced units such that particles have identical masses $m_i = 1$ and the temperature is chosen so that $\beta = 1$. The friction constant in the Langevin equations is $\gamma = 1$. Solvent particles interact with

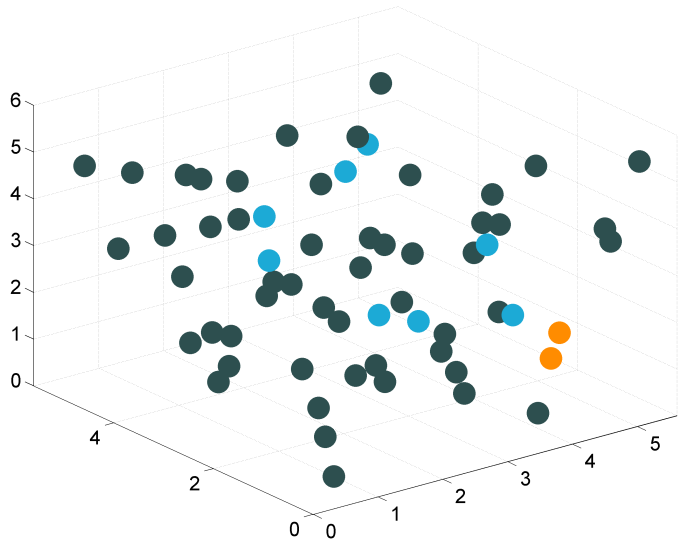

Figure 6: A snapshot from the AR-dynamics simulation of a dimer (orange) surrounded by 62 solvent particles (light blue and grey) with parameters $v_{\min} = 6.4$ and $v_{\max} = 8$ (only for the solvent particles, the dimer has the standard kinetic energy $k_{\mathrm{std}}(p_1)) + k_{\mathrm{std}}(p_2))$). The solvent particles are in two states: restrained (grey) or active (light blue).

each other and with the dimer particles by a truncated Lennard-Jones potential with parameter $\varepsilon_{\mathrm{LJ}} = 1$ with a cut-off distance $r_{\mathrm{LJ}} = 2^{1/6}$. Dimer particles interact with each-other with a double-well potential (with width $w = 1$ and height $h = 1$). This potential corresponds to a metastable system, with the two metastable states: compact and stretched. For a more precise formulation, see Section 5.2 in Redon et al. (2016) or Lelièvre et al. (2010). We discretize the modified Langevin equations (1) by a second-order scheme (OBABO) with time step size $\Delta t = 0.001$ and perform $N_{\mathrm{iter}} \approx 10^9$ time steps.

We use neighbor-lists based on the cut-off distance of the Lennard-Jones potential r_{LJ}, according to Algorithm 3. The average number of neighbors is estimated as $C = 0.25$. We run one reference simulation in the standard dynamics.

In the AR simulations, we consider non-zero restraining parameters on the solvent only, and we let dimer particles follow the standard dynamics. In order to demonstrate the dependence of the total speed-up S_{total} on the

restraining parameters $v_{\min}$ and $v_{\max}$, we consider the following observables: the dimer distance $A_D(q_1, q_2) = |q_1 - q_2|$, the dimer potential $A_V(q_1, q_2) = V_{DW}(|q_1 - q_2|)$ and the kinetic temperature $\mathcal{T}(p) = p \cdot \nabla K(p)$. The first two observables only depend on the positions of the dimer particles, hence we expect that the variance will not be much modified even for large restraining parameters. The function $\mathcal{T}(p)$ depends on the momenta of all particles p and satisfies $\langle \mathcal{T} \rangle_{\mu_{v_{\min}, v_{\max}}} = k_B T$. The knowledge of the exact average allows us to verify that the time step size Δt is chosen sufficiently small in order to make the systematic error on the averages smaller than 1% even for $v_{\max} = 10$. The asymptotic variance σ^2 of a time average for a given observable φ is estimated by approximating the integrated auto-correlation function by a trapezoidal rule (see Section 5.2 in Redon et al. (2016)).

First, we confirm theoretical predictions for the algorithmic speed-up S_{a}. In our simulations, we measure the time per force update, as well as the time per time step. We compare the measured speed-up, which is a ratio of the measured time in the standard dynamics and the AR dynamics, with the estimated speed-up in the force update (16) and for the overall time step (17). Figure 7 shows a comparison of the predicted algorithmic speed-up and the measured algorithmic speed-up in our simulation, which demonstrates that it is possible to roughly estimate the computational behavior of a specific implementation. Even for our simple implementation, however, the mismatch in the curves may have multiple causes: the update of positions and momenta, the creating of lists of active particles, the updating of the neighbor-lists *etc.*, and the fact that with growing $v_{\max}$ and a fixed δ, the particle spends more time in the transition region which is computationally more expensive due to the spline. As we suggest later in the paper, realistic implementations such as those found in popular MD packages (e.g. LAMMPS, GROMACS, etc., see Singh and Redon 2016) are very complex, and the best strategy to determine AR parameters may be to actually measure algorithmic speed-ups on short simulations.

Figure 8 plots the estimated relative variation of the variance of three observables as a function of the parameters $v_{\max}$ for $v_{\min} = \delta v_{\max}$ with $\delta = 0.5$. Recall that only the solvent particles are restrained and therefore the variance of an observable as $\mathcal{T}$, which depends directly on these degrees of freedom, is more perturbed from the variance in the standard case than the variance of an observable depending on particles following the standard dynamics (the dimer). We confirm the results showed in Redon et al. (2016):

the variance of $\mathcal{T}$ is modified more drastically than the variance of observables measured on the dimer with growing $v_{\max}$.

Finally, combining the algorithmic speed-up $S_{\mathrm{a}}(v_{\min}, v_{\max})$ with the variance $\sigma^2(v_{\min}, v_{\max})$, we estimate the total speed-up according to (9). This is depicted on Figure 9. We again consider $\delta_k \in \{0.5, 0.8, 0.9\}$ in order to demonstrate the impact of the gap between the parameters $\delta = v_{\min}/v_{\max}$ on the total speed-up S_{total}: the smaller the gap, the larger S_{total} becomes. Also, it holds that $S_{\mathrm{total}} > 1$ for the dimer observables only (up to 4), and not for the global observable $\mathcal{T}$, for which the relative deviation of the variance dominates the algorithmic speed-up. This supports the idea that we can speed-up the computation of macroscopic properties that depend on unrestrained degrees of freedom, i.e. those of the dimer in this example.

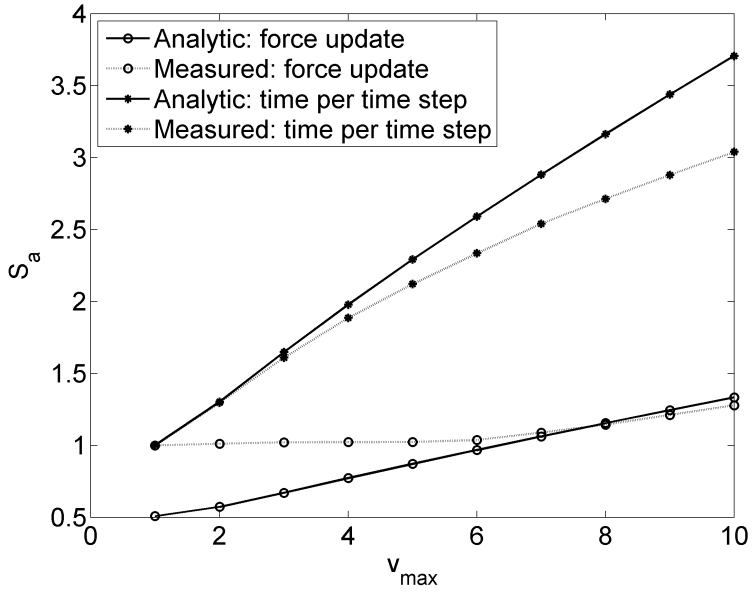

Figure 7: *Comparison of the analytically estimated algorithmic speed-up with the measured one.* We considered a system described in Section 6. We measured the algorithmic speed-up S_{a} of the forces update function and the total time step, with respect to the parameters $v_{\max}$ and $v_{\min} = 0.8 v_{\max}$. We observe the same behavior of the computed speed up in forces update (the one predicted by (16)) and the measured one, as well as the algorithmic speed-up per time step (17). Note that the measured S_{a} of the forces update is equal to one for small values of $v_{\max}$, which is due to the implementation of the condition on the adaptive forces update as proposed in Algorithm 3 which assures $S_{\mathrm{a}} \geq 1$.

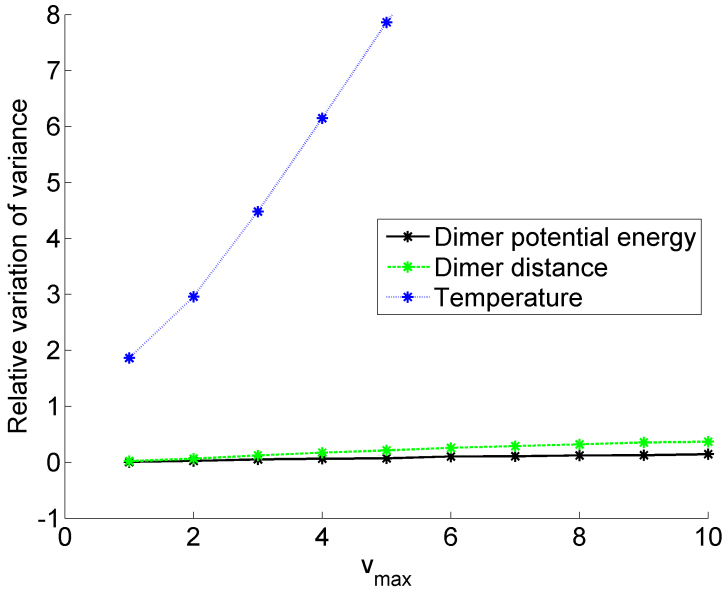

Figure 8: *Relative deviation of the asymptotic variance from the variance in the standard dynamics.* We considered a system of described in Section 6. We measured the variance of the dimer distance A_D (green), the dimer potential A_{DW} (black) and the temperature $\mathcal{T}$ (blue) in various parametrization of the AR dynamics. We plotted the relative deviation of the variance from the variance in the standard dynamics for parameters $v_{\max}$ with $\delta = 0.8$. Note that the variance is more perturbed with respect to the standard simulation for $\mathcal{T}$ than for the other two observables, which depend only on the dimer particles that are not restrained.

It is easy to compute the algorithmic speed-up S_{a}. The problematic part is the determination of the sensitivity of the observable on the restraining parameters (see again Figure 8). Since the variance can be approximated by a linear function of the restraining parameters at least locally, we can compute the slopes $c_{v_{\max}}$ such that[9] $\sigma^2(v_{\min}, v_{\max}) \approx \sigma^2(v_{\min}, v_{\max}) + c_{v_{\max}} v_{\max}$ from three AR simulations with parameters $(v^1_{\min}, v^1_{\max}), (v^1_{\min}, v^2_{\max}), (v^2_{\min}, v^1_{\max})$. More precisely, this allows us to approximate the total speed-up as

$$S_{\mathrm{total}}(v_{\min}, v_{\max}) \approx S_{\mathrm{a}}(v_{\min}, v_{\max}) \frac{1}{1 + \frac{c_{v_{\max}}}{\sigma^2(0,0)} v_{\max}} \,. \tag{23}$$

[9]Note that, in this linear approximation, we consider a fixed ratio δ such that $v_{\min} = \delta v_{\max}$.

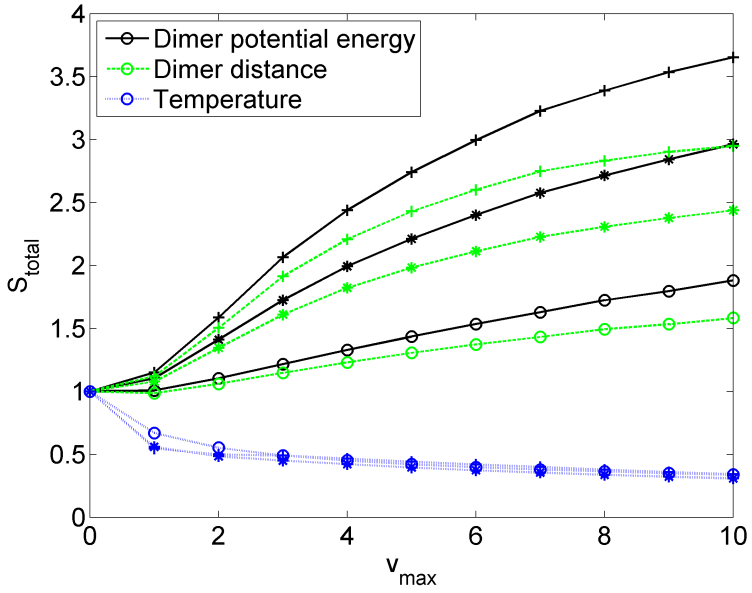

Figure 9: *Numerical estimation of the total speed-up depending on the parameters ratio.* We estimated the total speed-up S_{total} given by (9) for system described in Section 6. The algorithmic speed-up as well as the variance (see Figure 8) were obtained directly from the measurements in the simulation, *i.e.* not from the analytical formula. On the plot the results are showed for three observables (various color and line styles). The various markers correspond to different value of δ: the marker "open circles" corresponds to $\delta = 0.5$, the marker "star" corresponds to $\delta = 0.8$ and the marker "plus sign" corresponds to $\delta = 0.9$, such that $\delta = v_{\min}/v_{\max}$. The bigger is $v_{\min} = \delta v_{\max}$, the more particles are restrained which implies a bigger algorithmic speed-up. The bigger is δ, the less perturbed is the variance and hence the better is S_{total}.

We choose $(v_{\min}, v_{\max}) \in \{(3,6),(3,7),(2,6)\}$ and we estimate the slope $c_{v_{\max}}$. Table 1 shows the comparison with the slope directly obtained from simulations with fixed $\delta = 0.9$ and $\delta = 0.8$ for $v_{\max} \in [1, 10]$ (see Figure 8). This confirms that it is possible to capture the quantitative behavior of the relative slope $c_{v_{\max}}/\sigma_\varphi^2(0,0)$ from only three AR simulations. Note that the same approach could be used to determine the behavior of the variance as a function of different temperatures by measuring the variance only at a few points.

We have obtained an estimation of the variance $\sigma^2(v_{\min}, v_{\max})$. This allows us to express the total speed-up S_{total} as a function of $v_{\min}$ and $v_{\max}$, which is depicted on Figure 10 for A_D.

Remark 6.1. *It would be possible to push the parameters in Figure 10 in order to achieve a higher speed-up, up to the moment when the variance increase would start countering the algorithmic speed-up. Since the total speed-up depends on the simulated system and a concrete observable function, it does not make much sense to try and find the limit for our toy model. Moreover, for large parameters values, it is difficult to converge the quadratures in* (19) *and computationally too expensive to obtain the estimates numerically in the sense of Figure 10. Nevertheless, we believe Figure 10 provides a good understanding of the qualitative behavior of the total speed-up.*

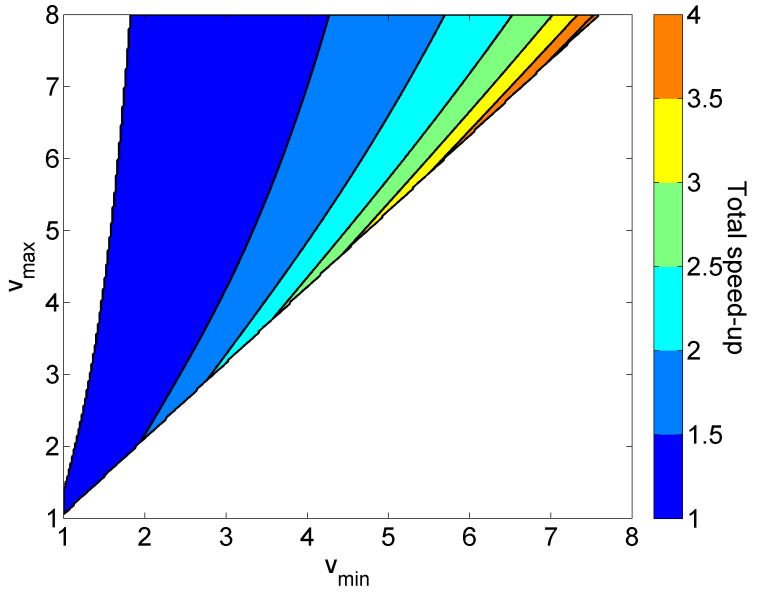

Figure 10: *Analytical estimation of the total speed-up.* We estimated the expected total speed-up S_{total} for the observable dimer distance A_D with respect to parameters $v_{\min}$ and $v_{\max}$ ($v_{\max} \leq 0.95 v_{\max}$). The variance was estimated from three points as a linear function of $v_{\min}$ and $v_{\max}$ and we used the analytical estimation of S_{a} according to (17). Only $S_{\text{total}} > 1$ is plotted.

7. Conclusions and future work

We have analyzed the speed-up achievable by AR Langevin dynamics in the estimation of thermodynamics properties by time averages, in particular

as a function of the restraining parameters $v_{\min}$ and $v_{\max}$. The final formula consists of two parts: the algorithmic speed-up and the modification of variance. The approach proposed in this work allows us to choose the parameters of the method. The theoretical results are confirmed by numerical experiments. We expect that even higher total speed-ups can be achieved when the complexity of the incremental force update algorithm is improved, for instance by avoiding the double calculation of inter-particle forces (in the add and subtract steps).

The stability of the AR dynamics can be rigorously analyzed and improved by introducing a Metropolis-Hasting step, which is the purpose of future work (see Stoltz and Trstanova (2016)). Furthermore, the AR dynamics can be extended to simulations under different temperatures, and hence combined with other methods, such as parallel replica exchange (Sugita and Okamoto (1999)). Another natural extension is to explore the dynamical properties of the modified dynamics, which is not obvious at first sight.

Acknowledgements

Stephane Redon and Zofia Trstanova gratefully acknowledge funding from the European Research Council through the ERC Starting Grant n. 307629. We also thank Gabriel Stoltz and Svetlana Artemova for helpful advice.

Bibliography

References

Allen, M. P., Tildesley, D. J., 1989. Computer simulation of liquids. Oxford University Press.

Artemova, S., Grudinin, S., Redon, S., 2011. A comparison of neighbor search algorithms for large rigid molecules. J. Comp. Chem. 32 (13), 2865–2877.

Artemova, S., Redon, S., 2012. Adaptively restrained particle simulations. Phys. Rev. Lett. 109 (19), 190201.

Balian, R., 2006. From microphysics to macrophysics: methods and applications of statistical physics. Vol. 2. Springer Science & Business Media.

Bennett, C. H., 1975. Mass tensor molecular dynamics. J. Comput. Phys. 19 (3), 267–279.

Bosson, M., Richard, C., Plet, A., Grudinin, S., Redon, S., 2012. Interactive quantum chemistry: A divide-and-conquer ASED–MO method. J. Comp. Chem. 33 (7), 779–790.

Espanol, P., Warren, P., 1995. Statistical mechanics of dissipative particle dynamics. Europhys. Lett. 30 (4), 191.

Fathi, M., Stoltz, G., 2015. Improving dynamical properties of stabilized discretizations of overdamped Langevin dynamics. arXiv preprint 1505.04905.

Frenkel, D., Smit, B., 2002. Understanding molecular simulation: from algorithms to applications. Elsevier.

Izaguirre, J. A., Hampton, S. S., 2004. Shadow hybrid Monte Carlo: an efficient propagator in phase space of macromolecules. J. Comput. Phys. 200 (2), 581–604.

Kopec, M., 2014. Weak backward error analysis for overdamped Langevin processes. IMA J. Numer. Anal. 55, 1057–1103.

Leimkuhler, B., Matthews, C., Stoltz, G., 2016. The computation of averages from equilibrium and nonequilibrium Langevin molecular dynamics. IMA J. Num. Anal. 36 (1), 13–79.

Lelièvre, T., Rousset, M., Stoltz, G., 2010. Free Energy Computations: A Mathematical Perspective. Imperial College Press.

Lelièvre, T., Stoltz, G., 2016. Partial differential equations and stochastic methods in molecular dynamics. Acta Numerica 25, 681–880.

Mattingly, J. C., Stuart, A. M., Higham, D. J., 2002. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. Stoch. Proc. Appl. 101 (2), 185–232.

Meyn, S. P., Tweedie, R. L., 1993. Stability of Markovian processes. II. Continuous-time processes and sampled chains. Adv. Appl. Probab. 25, 487–517.

Milstein, G. N., Tretyakov, M. V., 2013. Stochastic Numerics for Mathematical Physics. Springer Science & Business Media.

Plechac, P., Rousset, M., 2010. Implicit mass-matrix penalization of Hamiltonian dynamics with application to exact sampling of stiff systems. Multiscale Model. Sim. 8 (2), 498–539.

Redon, S., Stoltz, G., Trstanova, Z., 2016. Error analysis of modified Langevin dynamics. J. Stat. Phys. 164 (4), 735–771.

Singh, K. K., Redon, S., 2016. Adaptively restrained molecular dynamics in LAMMPS. Submitted to Comp. Phys. Comm.

Stoltz, G., Trstanova, Z., 2016. Stable and accurate schemes for Langevin dynamics with general kinetic energies. arXiv:1609.02891.

Sugita, Y., Okamoto, Y., 1999. Replica-exchange molecular dynamics method for protein folding. Chem. Phys. Lett. 314 (1), 141–151.

$c_{v_{\max}}/\sigma^2_{A_i}(0,0)$	$\delta = 0.9$: 3 points	$\delta = 0.9$: $v_{\max} \in [0, 10]$
A_{DW}	0.017863	0.016452
A_D	0.043855	0.042729
	$\delta = 0.8$: 3 points	$\delta = 0.8$: $v_{\max} \in [0, 10]$
A_{DW}	0.016426	0.015028
A_D	0.039587	0.039787

Table 1: Comparison of the relative deviation of the variance, obtained from linear interpolation of three points and the interpolation of values obtained from many simulations.